

SAM (REX/Ops) Operations Meeting 2007.10.02

Present: SteveW, SteveS, Randolph, Adam (typist), Art, Parag, Anil, Angela, Robert, Krzysztof, Valeria (remotely), Andrew, Dehong, Dave

For reports submitted to the meeting...

The report text is in this font

Comments made during the meeting are in this font inline with the report

Releases:

DB server testing of 8.3.0 (retired date) is still waiting on Randolph (will be made top priority this week).

Robert has made changes to DHGetLite - in CVS but not released yet.

Note "sam update dataset" documentation says you can add keywords for an existing dataset, but the server rejects these changes (since the first thing the server does is check if the dataset has any snapshots - if so, no changes are allowed). Robert's suggestion is to allow changes to description and keywords, so DB server needs to be changed (Adam adds to list). Since no one has noticed this problem until a few weeks ago, and this problem has been there since SAM's inception, use of this command is rare.

"sam list requests" command failure mystery is solved (see last week's report) - there was indeed an IDL change from DB server version 8.1 to 8.2. This change slipped through and makes old clients listing particular requests fail. Solution is to move to the new client (this is fine since only a very small number of people use this command). And the commandment is that in the rare case when an IDL struct needs to be changed, NEVER change it - add a new one instead to ensure backwards compatibility. The IDL that was changed was adding the number of files to the request handler table. Fortunately a minor change that only affects a small number of people.

Randolph wants to add an effective date to the work request table to be updated at the same time a request status is updated – this will make queries for dates much easier. Anil will add the column. Code will have to fill this column, and needs a script to populate the column for the first time (script is easy). Randolph will do the script and the coding.

DØ:

Kin Yip <kinyip@bnl.gov> Mon, Oct 1, 2007 at 4:32 PM

To: sam-design@fnal.gov

Events of concern/interest this week :

-- jetsam-clued0 and therefore "clued0" SAM station was down in the weekend and Robert rebooted it Monday (Oct. 1st).

As a shifter, I sent email to clued0-admin on Sat. and I understood that nothing might not happen during the weekend (though I was hoping for some miracles). There was only one ticket IT# 3054 complaining about this.

Robert rebooted jetsam and this fixed it. Robert will remind sam shifters of clued0 support policy (no one is going to come in on the weekend to reboot the machine). Shifters could contact the control room to have someone kick the machine if they can find it. This is the second crash recently - will use power outage to do diagnostics. Machine is still in warrantee.

-- fcpd problem on d0rsam01 and d0srv017 because fcpd is not on running on fnd0xxx machines such as fnd0621 .

Mike deemed that it's resolved after some configuration changes (IT #3046).

Not an fcpd problem on d0rsam01/d0srv017 - the problem was that Mike was trying to copy files from old farm nodes that were not running fcp (fcp must be running everywhere for transfers to work).

--- Complaints about "samgrid.fnal.gov not responding to web queries and MC requests" happened often.

condor_sched restarts automatically after crash - if try to access during restart, things won't work (crashed probably due to lots of condor_q commands - we should clean up the queue on samgrid.fnal.gov). Note this is an old version on condor_sched. Will upgrade with new VDT [virtual data toolkit] (Parag is still testing).

At the moment, SAM shifters do NOT have access to samgrid.fnal.gov and therefore can't help fix the problem, even though the procedure is clearly in FAQ # 81.

SAMGrid team and REX/ops will discuss giving access to samgrid.fnal.gov to shifters.

Unresolved issues (really no brand NEW and unresolved issues this week):

--- IT #3042 : Parag just created a ticket reminding ourselves that "Bug in Phased dataset jobs with notification in Samgrid JDF" is still being worked but not at the highest priority.

Since easy workaround, still not high priority.

--- tickets IT #3038/3030/3022 (getting UK files/xmlldb error at GridKa) all seem to be related to J. Snow.

See last week's report.

--- IT# 3023 : In the last meeting's minutes, it's said that "index" didn't solve the problem of slow response to the request `sam.getRequestInformation(afterDate=reqdategt)` . So, any news ? Query is very inefficient and is incorrect. We're having a meeting with production people to discuss solutions. Randolph has a prototype as a proof of principle of fast

and correct queries.

From Robert: Still having problems with extremely large log files on d0cabsrv1. Size is prop to the number of files being transferred. Robert discovered that one can just rename the log file and the station continues. Robert has a cron job that will mv the file when it gets too big (log file > 2 GB causes station crash and can't restart).

Todd Adams sent mail this morning asking if it was time to declare the end of SAM v5 MC. YEAH!! WooHoo!

CDF:

Valeria Bartsch <bartsch@fnal.gov> Tue, Oct 2, 2007 at 3:41 AM
To: Adam Lyon <lyon@fnal.gov>
Cc: sam-design@fnal.gov

Many dCache problems solved by Alex and Krzysztof.

Wednesday

=====

Alex restarted dcache pools on the nodes fcdldata136, 137, 155, 139 and 316. This was to cleanup a "cache allocate memory" error. Run2sys was needed to reload and restart of automounter to stop the error.

The pool to pool transfers at fcdldata111 seem to be stuck.

dCache pool fcdldata154 has memory allocation problems, same as fcdldata316 yesterday.
fcdldata150 is offline for upgrade.

Thursday

=====

Alex upgraded _some_ cdf dcache pools nodes to the recent code release, upgraded java to 64-bit version (where applicable) and restarted all pools in the TestPools group. Thus we may preview changes which will be deployed during upgrade.

Saturday

=====

One of the CDF enstore silos (LSM 0,1) was offline last night while the CE from StorageTek is diagnosing a problem with one of its two arms. The arm has been replaced within one hour. There was a row of tapes that, apparently, one arm can reach but the other can't. The CDF-9940B library was paused. Access to LTO3 tapes in the SL8500 at GCC is unaffected.
Thanks to enstore group for fixing this so quickly.

Monday

=====

Krzysztof moved 2 B Pools nodes to Test Pools and 2 ReadOnly Pools to B Pools. There was a problem with staging files in the B dataset dcache area. K. had forced a configuration reload for a portion of file servers and sees many files restore requests from tapes now.

There were a few more dCache cells offline, they were restarted.

MINOS:

Arthur Kreymer <kreymer@fnal.gov> Tue, Oct 2, 2007 at 8:50 AM
 To: sam-design@fnal.gov
 Cc: minos_sam_admin@fnal.gov

Minos status 2007 Oct 02

General -

Enstore -

New 'vault' raw data copies going to LTO-3
Probably going as I write this report. These are the safety copies (if all goes well, tapes are never read).

DCache -

SRM hung up Saturday night due to fndca2a monitor node crash
 Detected by an end user noting lack of FTP transfer web page
 Restarted Monday noon

fndca2a should be monitored by NGOP, generating alarm/page
 SRM is a critical service, should be monitored and supported
Why does SRM fail if a monitoring node goes offline? Because there's a critical Postgres database on this node (what's it for?). This machine was not 24x7 supported.

Will upgrade the support level. There was no training for support people about what was on this node and how to restart the services. And "cell services page" showed SRM as running when it was not. Support needs a test to check upness of SRM. Support has no way to check upness of SRM except for asking end user (Art). Dave will look into resolving these support issues.

Monitoring web pages were restarted (some down since 31 Jul)
WooHoo!

Slow ROOT DCache reads in root >= 5.12
We still need to rerun tests on larger test files.

BlueArc - Still awaiting deployment of SataBeast/BlueArc
as /minos/scratch and /minos/data
Will request mounting these on GP_GRID.
People needed to get this up are on vacation (back today).

SAM -

IT 1894 Need 'sam locate' in C++ API
Users really need this feature.

Oracle -

minosora3 Oracle upgrade to 10.2.0.3 today.
Request production upgrade before approx 23 Oct beam startup.
dev/int upgrade today. Production will be upgraded 10/18. Note that startup of MINOS data taking will start around 10/23, so please try to upgrade Oracle before then.

OPEN ISSUES

IT 1894: sam locate via C++ api, for simple client needs
IT 1979: Compiler warnings from sam_cpp_api
IT 2071: quadratic project startup time growth (from Nov 1 2005)

PLAN -

Migration to 9940B tape can proceed anytime - when ?
Update Minos dbservers and clients to v8

Art: From last week, mentioned that DB Browser is not on port 80 (a UK site security rules requires outside web servers to be on port 80). Randolph has put the servers on port 80 and is still going through security paperwork to have it seen from outside the lab.

DB:

Anil Kumar <akumar@fnal.gov> Tue, Oct 2, 2007 at 7:45 AM
To: sam-design@fnal.gov

D0 :

- Scheduling upgrade of d0ofint1/d0ofprd1 databases to 10.2.0.3 Oct 2, 2007
8:00AM-1:00PM

Anil says these upgrade are done successfully.

- Also Steve K. will shutdown the monitoring of the obsolete clariion array
and

remove some Clariion software packages.

Note: The array will still be attached and will not be removed until
November's downtime.

Adam suggests burial at sea for the old Clarion.

- SAM cut v6_10 on d0ofdev1 - To add composite index on lower(file_name) and
retired_date -done

- v6_4 index on create_date of data_files against d0ofdev1 done

- Service Request with Oracle on Full table scan Vs Index scan on
data_files_param_values.

Development is working on Bug Fix. Bug# 5601578 No time estimate yet for
bug fix.

SAR# 5837876.993

*Krzysztof: datafile_locations table emptied (again) in integration. Not sure why.
Randolph will look into it. Andrew says a similar problem happened with the D0
development database. There's no log per table, so can't tell when this happened.
Robert says it probably happened within 2.5 weeks since the CDF int station was still
up. Steve W says that the DB server has protection against deleting all data in a table.
Randolph has tried to put back some data into datafile_locations, but int station still
does not work. Perhaps a problem in data_storage_locations.*

*Why does this happen in integration and not in production? Note - if refresh from
production, autodev tables will get messed up. Will investigate.*

Enstore/dCache:

From Dave:

DØ is down now for enstore software updates. ADIC people are looking at problems with reading barcodes on one side of the robot. This problem accounts for why LTO2s go offline (can't read barcode). Failure rate on robot 2 is 10x more than on robot 1. Trying to convince ADIC that there is a real problem here (not just our regular problem with barcodes). All should be up by noon.

Also in process of removing DLTs. May remove them by hand (instead of ejecting them with the Robot). Problem with doing it by hand is that parts of tower are hard to reach and could be a longish downtime. Still deciding what to do.

CDF downtime tomorrow to upgrade dCache and pnfs server hardware (hopefully will improve performance and reliability). This work is scheduled from 8am-2pm.

Early Thursday morning, there is network maintenance that should be non-destructive. But if they have to reboot the switch, there will be a network glitch affecting enstore and dCache. But we think things should not be affected.

Next Thursday morning 10/11, there will be a network outage that will involve rebooting switch on FCC1. That will affect public dCache and public enstore. Stken will be down briefly for that time.

CDF dCache + silo problems (see above). DØ barcode problems (See above). Monitor node problems for SRM (see above). Public dCache running well since the upgrade (this upgrade will be applied to CDF tomorrow).